Backpressure on the Backbone: A Lightweight, Non-intrusive Traffic Engineering Approach

Christos Liaskos, Xenofontas Dimitropoulos and Leandros Tassiulas

arXiv:1611.05619v1 [cs.NI] 17 Nov 2016

***Abstract*—The present study proposes a novel collaborative traffic engineering scheme for networks of autonomous systems. Backpressure routing principles are used for deriving priority routing rules that optimally stabilize a network, while maximizing its throughput under latency considerations. The routing rules are deployed to the network following simple SDN principles. The proposed scheme requires minimal, infrequent interaction with a central controller, limiting its imposed workload. Furthermore, it respects the internal structure of the autonomous systems and their existing peering relations. In addition, it co-exists smoothly with underlying distance vector-based routing schemes. The proposed scheme combines simplicity with substantial gains in served transit traffic volume, as shown by simulations in realistic setups and proven via mathematical analysis.**

***Index Terms*—Traffic engineering, autonomous systems, backpressure routing.**

I. Introduction

SOFTWARE-Defined Networking (SDN) can imbue the network management process with an unparalleled level of state monitoring and control. The ability to migrate the routing elements of a network from closed, static hardware solutions towards an open, re-programmable paradigm is expected to promote significantly the adaptivity to time-variant demand patterns, eventually yielding a healthy and constant innovation rate. The OpenFlow protocol and assorted hardware [1], which enables an administrative authority to centrally monitor a network and deploy fitting routing strategies, has already produced significant gains in a wide set of application scenarios [2], [3].

Traffic engineering (TE) constitutes one of the networking aspects that has benefited the most from the adoption of centralized control. The objective of TE is the real-time grooming of data flows, in order to provide the best possible quality of service on a given physical infrastructure. Therefore, the success of TE depends on the ability to obtain accurate snapshots of the network's state in little time, as well as to react rapidly to state changes [4]. Given that SDN offers a unique advantage in terms of centralized network monitoring and fine-grained actuation, it is not surprising that the TE and SDN combination has been highly fruitful [5]. Early industrial applications have yielded considerable gains in the efficient use of network resources, such as link utilization [6] and application-specific network service [2].

Given the success of SDN solutions in proprietary networks such as datacenters [2], [6], research aims to bring its benefits to the Internet routing [7]. Current studies promise the provision of end-to-end QoS guarantees at Internet-level [8], [9], fast and secure convergence of global TE [10], and enabling the continuous evolution of Internet routing [11]. Nonetheless, adopting such solutions at global level faces the challenge of multi-domain coordination. The Internet comprises Autonomous Systems (ASes), which in their majority are privately held and managed organizations. On one hand, each AS may have its own hardware, internal network control system, as well as financial status and business model. On the other hand, related solutions promise impressive performance but typically require changes in the AS internal state of affairs, at least partially [8]–[10]. Thus, while collaborative TE may be to the financial interest of several ASes [12], [13], more intermediate steps may be required for promoting it.

The present study contributes a light collaborative TE solution for AS collaboration, where participants maintain full control of their own networks. The methodology consists of applying the principles of Backpressure (BPR) routing to a network of ASes [14]. BPR is well-known for its simple principles, leading to analytically-proven throughput-optimality nonetheless. Within the proposed system, collaborating ASes inform a central service of internal congestion events that they experience. This information can be derived in any manner the AS presently supports. Moreover, it refers to the destination of the congested traffic only and, therefore, is partial and of temporal validity. In return, congested ASes receive a small set of proposed priority routing rules derived by BPR. The core BPR principle is the offloading of the congested traffic to the least loaded neighboring AS. ASes can decide upon the application of the proposed rules and proceed to their deployment in any custom manner. No further details on the inner workings of an AS need to be shared. The proposed scheme can respect existing routing algorithms, peering preferences and latency considerations, while offering analytically-proven throughput maximization and network stabilization potential. Moreover, throughput maximization can translate to substantially increased transit traffic, which can be of financial interest to the collaborating ASes. The proposed scheme is intended as an intermediate step towards fostering AS collaboration. Its aim is to serve as a stepping stone towards more

C. Liaskos is with the Foundation for Research and Technology Hellas (FORTH), Institute of Computer Science, N. Plastira 100, Vassilika Vouton, GR-700 13 Heraklion, Crete, Greece. e-mail: cliaskos@ics.forth.gr.

X. Dimitropoulos is with the Computer Science Department, University of Crete, and the Foundation for Research and Technology Hellas (FORTH), Institute of Computer Science, N. Plastira 100, Vassilika Vouton, GR-700 13 Heraklion, Crete, Greece. e-mails: fontas@csd.uoc.gr, fontas@ics.forth.gr.

L. Tassiulas is with the School of Engineering and Applied Science, Yale University, P.O. Box 208263, New Haven, CT 06520, e-mail: leandros.tassiulas@yale.edu.

This work was funded by the European Research Council, grant EU338402, project NetVOLUTION (http://www.netvolution.eu).

Table I: Summary of Notation

Symbol	Explanation
$U_{(n,c)}(t)$	The aggregate traffic accumulated within a network node n at time t, destined towards node c.
T	The network state monitoring and actuation period.
$O_{(n,c)}^{t\to t+T}$	Data volume outgoing from node n to c in the interval $[t,\, t+T]$.
$I_{(n,c)}^{t\to t+T}$	Data volume incoming to node n, destined to c, in the interval $[t,\, t+T]$.
$G_{(n,c)}^{t\to t+T}$	Data volume generated at node n, (or arriving to n from a network-external source), destined to c, in the interval $[t,\, t+T]$.
$\mu_l^{(c)}(t)$	The maximum allowed bitrate at time t over a network link l, carrying traffic destined to node c. For wired backbone links, $\mu_l^{(c)}$ is the nominal capacity of l, and is invariant of t.
$\overline{\lambda_{(n,c)}}(t)$	The average traffic production rate at node n destined to c for the time interval $[t, t+T]$.
$n' \leftarrow \mathcal{P}(n)$	The application of a routing policy $\mathcal{P}$ on data residing at node n, returning the next intended hop n'.
$\overrightarrow{T}\{\mathcal{P}^{\circ m}(n)\}$	The set of traversed nodes when applying m times the policy $\mathcal{P}$ on node n.

Algorithm 1 The Standard Backpressure Routing algorithm.

1: **procedure** SBPR($network_state(t = mT | m \in \mathbf{N})$)
2: **for** each node c **do** ▷ Assign traffic type to links.
3: **for** each link l **do**
4: $c_l^*(t) \leftarrow argmax_c\{U_{source(l),c} - U_{dest(l),c}\}$
5: $\Delta Q_l^*(t) \leftarrow max_c\{0, U_{source(l),c} - U_{dest(l),c}\}$
6: **end for**
7: **end for** ▷ Calculate optimal transfer rates per link.
8: $\mu^*(t) \leftarrow argmax_\mu \sum_{\forall l} \mu_l(t) \cdot \Delta Q_l^*(t)$
9: **for** each link $l : \Delta Q_l^*(t) > 0$ **do**
10: Transfer data destined to $c_l^*(t)$ with rate $\mu_l^*(t)$.
11: **end for**
12: **end procedure**

advanced SDN schemes for AS collaboration [8], [9], [11]. It prioritizes compatibility with existing AS equipment and lays a foundation for centralized, SDN-based inter-AS routing orchestration.

The remainder of this paper is as follows. Prerequisites are given in Section II. The system model and the analysis follow in Sections III and IV. Evaluation via simulations takes place in Section V. Section VI presents the related studies and the paper is concluded in Section VII.

II. Prerequisites

An important term in networking studies is the notion of *network stability*. It is defined as the ability of a routing policy to keep all network queues bounded, provided that the input load is within the network's traffic dispatch ability, i.e. within its *stability region*. Let $U_{(n,c)}(t)$ denote the aggregate traffic accumulated within a network node n at time t, destined towards node c. Stability is then defined as [15, p. 24]:

$$\limsup_{\tau\to+\infty} \frac{1}{\tau} \sum_{t=1}^{\tau} E\left\{U_{(n,c)}(t)\right\} < \infty,\ \forall n, c \tag{1}$$

where τ is the time horizon and $E\{*\}$ denotes averaging over any probabilistic factors present in the system.

The Standard BPR routing algorithm (SBPR, Algorithm 1) has been proven to optimally stabilize a network, i.e., keeping it within its stability region and maximizing its throughput [16], [17]. Its operating principle is simple: a node n offloads traffic towards a destination c by redirecting it to its less congested, immediate neighbor (lines $2-7$). If the interconnecting links have non-constant bandwidth (i.e., wireless links), a selection of valid rates takes place in line 8, and the data is transferred in lines $9-11$.

The SBPR algorithm was derived analytically, using a well-known network stability framework, i.e., the Lyapunov Drift. This approach defines a quadratic function of the form:

$$L(t) = \sum_n \sum_c U_{(n,c)}^2(t) \tag{2}$$

The goal is then to deduce the bounds of $\Delta L(t) = E\left\{L(t+T) - L(t)\right\}$, which describes the evolution of the network queue levels over a time period T. The *Lyapunov stability theorem* states that if it holds [15, p. 50]:

$$\Delta L(t) \le B - \epsilon \cdot \sum_n \sum_c U_{(n,c)}(t) \tag{3}$$

for any two positive quantities, B and ϵ, then the network is stable and average queue size of inequality (1) is bounded by B/ϵ instead of drifting towards infinity.

The *BPR* class of algorithms defines a routing policy that complies with the stability criteria of inequality (3). Its goal is to minimize the lower bound of $\Delta L(t)\,\forall t$, effectively suppressing the average queue level within the network. The analytical approach, followed by related studies [18], [19], is outlined as follows.

Firstly, the queue dynamics are expressed by the following relation, $\forall n, c$ [15, p. 54]:

$$U_{(n,c)}^{(t+T)} = max\left\{0, U_{(n,c)}^{(t)} - O_{(n,c)}^{t\to t+T}\right\} + I_{(n,c)}^{t\to t+T} + G_{(n,c)}^{t\to t+T} \tag{4}$$

where $O_{(n,c)}^{t\to t+T}$, $I_{(n,c)}^{t\to t+T}$ and $G_{(n,c)}^{t\to t+T}$ denotes outgoing, incoming and locally generated data, as summarized in Table I. Notice that these quantities refer to any traffic exchange process in general, and not just to traffic exchanged due to BPR routing.

Secondly, both sides of equation (4) are squared, and a series of relaxations are applied on its right hand side (RHS). These relaxations are based on the identity [15, p. 53]:

$$V \le max\{0, U-\mu\} + A \Rightarrow V^2 \le U^2 + \mu^2 + A^2 - 2U\cdot(\mu - A) \tag{5}$$

as well as on the following inequalities, $\forall n, c$:

$$O_{(n,c)}^{t\to t+T} \le \sum_{l:\, source(l)=n} \int_t^{t+T} \mu_l^{(c)}(t)\cdot dt \tag{6}$$

$$I_{(n,c)}^{t\to t+T} \le \sum_{l:\, destination(l)=n} \int_t^{t+T} \mu_l^{(c)}(t)\cdot dt \tag{7}$$

where $\mu_l^{(c)}(t)$ is the available capacity of a network link l at time t, that can carry traffic destined to node c.

Thus, one derives an inequality of the form of relation (3). Further relaxation by substituting all $\mu_l^{(c)}$ and $G_{(n,c)}^{t\to t+T}$ with

maximum, nominal values yields compliance with the Lyapunov stability theorem. Furthermore, it is deduced that the upper bound of relation (3) can be minimized when maximizing the quantity [15, p. 66]:

$$\sum_{\forall n,k,c} \mu^{(c)}_{l:source(n)\rightarrow dest(k)}(t) \cdot \left(U_{(n,c)}(t) - U_{(k,c)}(t)\right) \quad (8)$$

The SBPR algorithm (Algorithm 1) essentially expresses the optimization pursuit of relation (8) at lines $2-7$.

The BPR class of algorithms has found extensive use in packet switching hardware, wireless ad hoc networks and satellite systems due to their throughput optimality trait [17]–[19]. However, if used as a standalone routing policy with no other support, BPR can yield increased latency. This is due to the fact that queue sizes must build up in order to obtain backpressure potential, ΔU, yielding increased queuing delay. On the other hand, when the queue sizes are low, routing may resemble a random walk, accentuating propagation delays. Furthermore, if T is smaller than the full trip time of a packet, loops may appear in its path. Newer algorithms of the BPR class mitigate these issues and take into account latency considerations. For example, authors in [20] restrict steps $3-5$ of Algorithm 1 only within a subset of links that offer a bounded maximum number of hops towards the target. Other studies have shown that simply altering the queuing discipline from FIFO to LIFO yields considerable latency gains [21]–[23]. Finally, it is worth noting that the BPR class can be made TCP compatible in a straightforward fashion [24].

Another major advantage of BPR is the fair distribution of traffic over the nodes. Particularly, it has been shown that the closer a node is to the gateway(s), the more the transit traffic will pass through it over the time interval T [23]. Thus, nodes surrounding a gateway will naturally receive more transit traffic than a distant node. Nonetheless: i) nodes with similar distance and connectivity to the gateway(s) will receive similar transit traffic volume, while ii) all nodes in general will receive surplus traffic [23]. In addition, the variance of the transit traffic served by each node is minimal and very close to their long-term average [22]. Therefore, BPR shares the transit traffic fairly among the network nodes, with regard to their connectivity and their placement within the topology.

III. The System Model

The present paper proposes the use of BPR policies in backbone networks. The assumed setup, given in Fig. 1, considers a network comprising border routers organized in ASes (nodes). These ASes constitute a cluster of collaborating, same-tier, transit service providers (ISPs). The ASes exchange traffic with the rest of the Internet normally (i.e., they are connected to other, non-collaborating ASes) via arbitrary links. The proposed scheme seeks to promote the creation of such collaboration clusters, by yielding a significant increase in the total transit traffic serviced by the cluster.

The proposed scheme prioritizes minimal operational and equipment changes in the AS cluster. Towards this end: i) each AS is allowed to retain its existing Network Control System

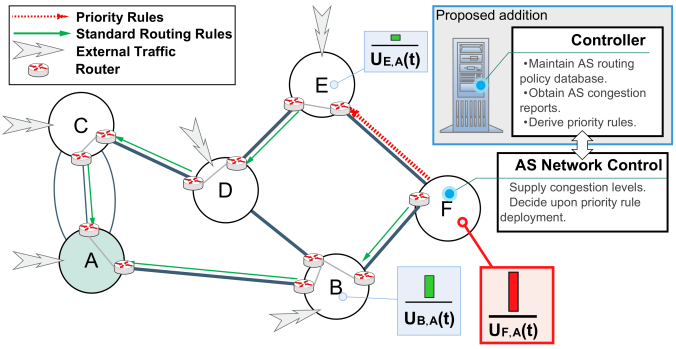

Figure 1: The employed system setup. A network of ASes, A-F, uses BPR-derived routing rules on top of its standard routing scheme. A centralized control entity supplies the temporary priority rules on demand, while existing AS Network Control Systems retain their monitoring and routing jurisdiction.

(NCS), which comprises AS-local router-level topology maintenance, congestion measurement services, link quality monitoring and routing rule management. The proposed scheme essentially operates on top of any set of NCSes (legacy or SDN-based, such as ONOS [25] or Beehive [26]), making use of their existing functionality. ii) The cluster of ASes may have any load-invariant routing policy, e.g., Distance Vector Routing (DVR) enforced via the Border Gateway Protocol (i/eBGP) [27]. The proposed scheme assumes no changes or obstruction to the BGP operation. ASes within the collaboration cluster exchange BGP updates with other ASes normally, whether these are members of the collaboration cluster or not. Routing policies (i.e., AS routing preferences) are allowed, as per the existing status quo in inter-AS routing. iii) The proposed scheme functions by installing priority routing rules at *adjacent* border routers, essentially operating on top of the underlying BGP routing. These priority rules are of *temporal* value and aim at alleviating AS-local increases in traffic load. As such, the priority rules *do not* propagate in the form of updates, either within or outside the cluster.

The proposed scheme considers the addition of a central Controller, which supplies each AS-local NCS with the proposed routing rules upon demand. An NCS is assumed to monitor its internal congestion level using its existing components. On the event of increased congestion (or periodically), it forwards the congestion measurement to the Controller, obtaining a set of proposed priority routing rules in return. The NCS then decides upon their deployment or rejection. Exemplary reasons for rejection are a sudden change in the AS routing policies or reference to faulty/flapping links, as detected by the NCS. Subsequently, the Controller is updated on the future AS policy and set of usable links.

The traffic volume metric monitored by the NCS of each AS, n, is essentially $U_{(n,c)}(t)$ (cf. Table I). Therefore, we notice that the congestion information passed to the Controller is partial and specific to the congestion event. Obtaining $U_{(n,c)}(t)$ is handled by an NCS, using any existing module or approach. For instance, the *router-state polling* technique measures the ingress and egress traffic volume that has traversed the interfaces of the AS border routers over a period of time. The quantity $U_{(n,c)}(t)$ is then calculated simply subtracting the

total egress from the total ingress load, for every AS $c \neq n$ in the cluster. This type of monitoring can be accomplished in a very short amount of time ($1-5$ sec), even for very large networks [28]. Without loss of generality, we will assume that the Controller is supplied with the $U_{(n,c)}(t)$ measurements with period T, for every n and any destination c that exhibits congested traffic within n.

The Controller also acts as a repository holding the inter-AS routing preferences within the collaboration cluster. Each collaborating AS is assumed to regularly submit and update its routing preferences to this repository. The submitted information remains private and is used only internally by the Controller to derive proper priority rules. The collaborating ASes may also inform the controller of their topology, i.e., their inter-AS *usable* links and corresponding capacities. Notice this information can be presently approximated from public databases as well [29]. Having the routing preferences, the topology and the $U_{(n,c)}(t)$ information at its disposal, the Controller executes a BPR algorithm (e.g., Algorithm 1) at AS-level.

An example is shown in Fig. 1, where the Controller seeks to offload the traffic volume $U_{(F,A)}(t)$ at node F. A BPR algorithm is executed, which deduces that traffic from F towards A should better be offloaded to neighboring node E for the time being. A corresponding priority routing proposal is passed to AS F, which handles its deployment. The installed priority rule takes precedence over all other routing rules pertaining to link l_{FE}. Other routers within AS F (e.g., using iBGP) are then instructed to forward traffic for AS A towards the selected border router(s) by the active NCS.

Peering agreements and routing preferences are handled during the execution of the BPR algorithm. For example, returning to the example of Fig. 1, the controller would not propose the illustrated flow rule if it was disallowed by the respective peering preferences/agreements affecting F and E. In other words, when the BPR-variant searches for neighbors $s \in S : \{U_{(s,A)}(t) < U_{(F,A)}(t)\}$, the search is assumed to be limited to nodes that comply with any form of preference of agreement.

The management of the derived priority rules follows the general principles of the OpenFlow protocol [1]. Their lifetime is ended by the AS when the local congestion drops beyond a safety-level, or if the rule has remained idle for a period of time. Thus, the AS system reverts to its DVR routing policy as soon as possible. The Controller participates in the management of the rules as well. Similarly to the OpenFlow case, basic concerns when installing the priority routing rules are: i) the avoidance of routing loops, and ii) the protection against partial or asynchronous rule deployment. In this aspect, the Controller can implement existing algorithms for the priority rule derivation, which promise scalability and disruption-free network operation [30]. An interesting trait of the proposed scheme is that it can guarantee a basic, loop-free operation, even without such a rule management mechanism. This trait is based on a smart filtering of AS-neighbors when applying BPR and is analyzed in Section IV.

It is also worth noting the motives for participation of an AS to the proposed collaboration scheme, which derive from the inherent properties of BPR:

- Due to throughput optimality, the nodes are expected to handle more transit traffic, i.e., incoming to the cluster of collaborating nodes, using the same links, postponing the need for capacity upgrades.
- The incoming transit traffic is naturally shared fairly among the collaborating nodes, considering their connectivity and position in the topology, as discussed in Section II. This can be an attractive trait in the cases where the monetary profit of the ASes is strongly related to the served traffic [13]. In such cases, the pricing and profit model can remain unaltered, without necessarily requiring a profit-sharing mechanism.
- Natural traffic sharing and throughput optimality may also imply increased resilience against link-outages, attributed to benign causes (e.g., link failures, BGP session resets, link flooding due to flash-crowds [31]) or attacks (e.g., [32], [33]). The latter have recently shown potential to even cut-off ASes from the Internet [34]. In this case, collaboration can potentially postpone the need for sophisticated-yet very intrusive-defense mechanisms [35].
- The nodes remain autonomous, and need not reveal their internal structure/equipment/client set or relinquish control over it.
- BPR relies on node-local traffic state information only (cf. Algorithm 1). The required information is a common congestion metric.
- Controller failures are not critical, as on such events the ASes continue their normal, BGP-based operation.

In summary, the key-benefit of the proposed scheme is that it promises increased stability and transit traffic, while incurring minimal changes in the operation of the ASes. The proposed scheme follows some general OpenFlow operation principles, namely the central derivation of priority flow rules, but does not require upgrades to OpenFlow equipment. To the contrary, it is designed to work with the existing NCS that an AS has adopted. The main implementation cost for the proposed scheme is thus that of the Controller. In this aspect, the operation that the Controller should support does not differ significantly from a web service, which constitutes an indication of limited capital and operational expenditure.

IV. Analysis

The target of the analysis is to evaluate the use of BPR routing in backbone networks. To this end, the following sub-objectives are studied separately:

- The potential of BPR algorithms to take future traffic states into account.
- The system's dependence on the controller. To this end, we study the effects of the monitoring and actuation period T on the network drift bound.
- The loop-free, latency-aware cooperation between BPR and the underlying routing scheme, considering asynchronous or partial deployments of BPR rules.

Finally, we allow for at most a single priority flow rule per physical network link, which will be shown to be sufficient for significant performance gains.

A. On traffic forecast-awareness in BPR routing.

SBPR can serve as the BPR algorithm in the proposed system model of Section III. SBPR operates without knowledge of future traffic patterns, based on current congestion states only. Nonetheless, backbone links are characterized by high capacity. At such nominal data rates, the congestion distribution of ASes could vary rapidly. Moreover, ISPs already employ Internet traffic forecasts to their present TE approaches [36], [37]. These forecasts are quite accurate, indicatively yielding $1-3$ % error for $T=1\rightarrow 5$ min and $3-5$ % for $T=1$ hour [38]. Thus, the posed question is whether a new BPR algorithm can consider this additional information, without altering the properties of SBPR.

We begin the analysis by simplifying the RHS of relation (7), based on the fact that the network links have time-invariant bandwidth:

$$I_{(n,c)}^{t\rightarrow t+T}\leq \sum_{l:\, d(l)=n}\int_{t}^{t+T}\mu_{l}^{(c)}dt=T\cdot\sum_{l:\, d(l)=n}\mu_{l}^{(c)} \quad (9)$$

In the same manner, the RHS of relation (6) becomes:

$$O_{(n,c)}^{t\rightarrow t+T}\leq \sum_{l:\, s(l)=n}\int_{t}^{t+T}\mu_{l}^{(c)}dt=T\cdot\sum_{l:\, s(l)=n}\mu_{l_{nb(n)}}^{(c)} \quad (10)$$

where $b(n)$ represents the neighboring node of n, at the end of each outgoing link. The considered links are the ones compliant with any bilateral routing preferences. Moreover, for ease of presentation we set:

$$\mathring{\mu}_{(n)}^{(c)}=\sum_{l:\, s(l)=n}\mu_{l_{nb(n)}}^{(c)} \quad (11)$$

We proceed to apply identity (5) to eq. (4):

$$U_{(n,c)}^{2}(t+T)\leq U_{(n,c)}^{2}(t)+\left[O_{(n,c)}^{t\rightarrow t+T}\right]^{2}+\left[I_{(n,c)}^{t\rightarrow t+T}+G_{(n,c)}^{t\rightarrow t+T}\right]^{2} - 2\cdot U_{(n,c)}(t)\cdot\left[O_{(n,c)}^{t\rightarrow t+T}-I_{(n,c)}^{t\rightarrow t+T}-G_{(n,c)}^{t\rightarrow t+T}\right] \quad (12)$$

Using the updated relaxations (9) and (10) and setting $\Delta U_{(n,c)}^{2}(t)=U_{(n,c)}^{2}(t+T)-U_{(n,c)}^{2}(t)$ for brevity:

$$\Delta U_{(n,c)}^{2}(t)\leq T^{2}\cdot\left(\mathring{\mu}_{(n)}^{(c)}\right)^{2}+\left[T\cdot\sum_{l:\, d(l)=n}\mu_{l}^{(c)}+G_{(n,c)}^{t\rightarrow t+T}\right]^{2} - 2\cdot U_{(n,c)}(t)\cdot\left[T\cdot\mathring{\mu}_{(n)}^{(c)}-T\cdot\sum_{l:\, d(l)=n}\mu_{l}^{(c)}-G_{(n,c)}^{t\rightarrow t+T}\right] \quad (13)$$

It is not difficult to show that the RHS of inequality (13) can be reorganized as:

$$\Delta U_{(n,c)}^{2}(t)\leq\left[T\cdot\sum_{l:\, d(l)=n}\mu_{l}^{(c)}+U_{(n,c)}(t)+G_{(n,c)}^{t\rightarrow t+T}\right]^{2} + \left[T\cdot\mathring{\mu}_{(n)}^{(c)}-U_{(n,c)}(t)\right]^{2}-2\cdot U_{(n,c)}^{2}(t) \quad (14)$$

Summing both sides $\forall n,c$ and reminding that $\Delta L(t)=\sum_{\forall n}\sum_{\forall c}\Delta U_{(n,c)}^{2}(t)$:

$$\Delta L(t)\leq\sum_{\forall n}\sum_{\forall c}\underbrace{\left[T\cdot\sum_{l:\, d(l)=n}\mu_{l}^{(c)}+U_{(n,c)}(t)+G_{(n,c)}^{t\rightarrow t+T}\right]^{2}}_{(B)} + \sum_{\forall n}\sum_{\forall c}\underbrace{\left[T\cdot\mathring{\mu}_{(n)}^{(c)}-U_{(n,c)}(t)\right]^{2}}_{(A)}-2\cdot\sum_{\forall n}\sum_{\forall c}U_{(n,c)}^{2}(t) \quad (15)$$

Relation (15) provides some early insights on the minimization of the upper bound of the Lyapunov drift, $\Delta L(t)$. First of all, the terms (A) and (B) comprise sums of squares, which can be minimized by ideally nullifying each squared quantity. Term (A) advocates for long-lived routing decisions and against bandwidth waste. Furthermore, outflows from a term (A) will be added to a term within (B). While the present load of the recipient node, $U_{(n^{*},c)}(t)$ could be low or null, it may be expected to, e.g., generate a considerable load locally in the time interval $t\rightarrow t+T$ (expressed via $G_{(n,c)}^{t\rightarrow t+T}$). To quantify this concerns, we treat the RHS of relation (15) as a function of the BPR-derived routing decisions $\mathring{\mu}_{(n)}^{(c)}$ and attempt a straightforward optimization. The $\mathring{\mu}_{(n)}^{(c)}$ can be initially treated as continuous variables. Once optimal values have been derived, they can be mapped to the closest of the actually available options within the network topology. The sufficient conditions for the presence of a minimum are:

$$\begin{cases}\frac{\partial RHS_{(15)}}{\partial\mathring{\mu}_{(n)}^{(c)}}=0 & (i)\\ 0<\mathcal{H}\left(\frac{\partial^{2}RHS_{(15)}}{\partial\mathring{\mu}_{(n)}^{(c)}\cdot\partial\mathring{\mu}_{(k)}^{(c)}}\right)<\infty & (ii)\\ RHS_{(15)}\text{ is convex w.r.t. }\mathring{\mu}_{(n)}^{(c)},\mathring{\mu}_{(k)}^{(c)} & (iii)\end{cases} \quad (16)$$

where k denotes a node, $\mathcal{H}$ is the Hessian matrix [39] and the requirement $0<\mathcal{H}<\infty$ refers to each of its elements. From condition (16-i) we obtain:

$$2T\cdot\left[T\cdot\sum_{l:\, d(l)=b(n)}\mu_{l}^{(c)}+U_{(b(n),c)}(t)+G_{(b(n),c)}^{t\rightarrow t+T}\right] + 2\cdot T\left[T\cdot\mathring{\mu}_{(n)}^{(c)}-U_{(n,c)}(t)\right]=0\Longleftrightarrow \quad (17)$$

$$T\cdot\left[\sum_{l:\, d(l)=b(n)}\mu_{l}^{(c)}+\mathring{\mu}_{(n)}^{(c)}\right] - \left[U_{(n,c)}(t)-\left(U_{(b(n),c)}(t)+G_{(b(n),c)}^{t\rightarrow t+T}\right)\right]=0,\ \forall n,c \quad (18)$$

It is not difficult to show that conditions (16-ii, iii) are satisfied, since it holds:

$$\frac{\partial^{2}RHS_{(15)}}{\partial\mathring{\mu}_{(n)}^{(c)}\cdot\partial\mathring{\mu}_{(k)}^{(c)}}\propto T^{2}>0,\ \forall n,k \quad (19)$$

Equation (18) represents a generalization over the SBPR (Algorithm 1). At first, eq. (18) defines a linear system with discrete variables $\mathring{\mu}_{(n)}^{(c)}$ and can be solved as such. However, interesting approximations can be derived, which also exhibit

Algorithm 2 The Foresight-enabled Backpressure Routing algorithm.

1: **procedure** FBPR($network_state(t = mT | m \in \mathbf{N})$)
2: **for** each node n **do** ▷ Define priority flows.
3: $visited[c] \leftarrow 0, \forall c$
4: **for** each link $l : source(l) = n$ **do**
5: $c_l^*(t) \leftarrow \underset{c:!visited[c]}{argmax} \{U_{(n,c)} - U_{(d(l),c)} - G_{(d(l),c)}^{t \to t+T}\}$
6: $visited[c_l^*(t)] \leftarrow 1$
7: $\Delta Q_l^*(t) \leftarrow max\{0, U_{(n,c_l^*(t))} - U_{(d(l),c_l^*(t))}\}$
8: **end for**
9: **end for** ▷ Consider multi-links, if any.
10: $\mu^*(t) \leftarrow argmax_\mu \sum_{\forall l} \mu_l \cdot \Delta Q_l^*(t)$
11: **for** each link $l : \Delta Q_l^*(t) > 0$ **do**
12: Deploy rule $\{from : s(l), to : c_l^*(t), via : l\}$.
13: **end for**
14: **end procedure**

a dependence of the optimal solution from the network traffic forecasts.

Firstly, we notice that the term $T \cdot \left[\sum_{l:\, d(l)=b(n)} \mu_l^{(c)} + \mathring{\mu}_{(n)}^{(c)} \right]$ includes *nominal* link capacities μ. Therefore, it represents the aggregate traffic destined to c that could reach node $b(n)$, if all links were used exclusively and concurrently for this task, using their full capacity for the time interval T. In other words, the term is generally an upper bound of the traffic incoming to $b(n)$. Therefore, the focus is on the case where:

$$\sum_{l:\, d(l)=b(n)} T \left(\mu_l^{(c)} + \mathring{\mu}_{(n)}^{(c)} \right) > \underset{(n,c)}{U}(t) - \left(\underset{(b(n),c)}{U}(t) + G_{(b(n),c)}^{t \to t+T} \right) \tag{20}$$

In this case, the equality in equation (18) cannot be enforced. However, $RHS_{(15)}$ is convex, as shown in eq. (19). Therefore, $RHS_{(15)}$ is minimized when its first derivative, i.e., equation (18), is closest to zero. This is obtained by maximizing:

$$\Delta_{(n,c)}(t) = \left[U_{(n,c)}(t) - \left(U_{(b(n),c)}(t) + G_{(b(n),c)}^{t \to t+T} \right) \right] \tag{21}$$

which depends on the traffic generated locally at node $n(b)$ during $[t, t+T]$. In other words, the throughput-optimizing routing decision at node n, regarding traffic destined to node c are derived as follows:

$$n^* = argmax_{b(n)} \left\{ \Delta_{(n,c)}(t) \right\} \tag{22}$$

where n^* is the optimal neighboring of n to offload data to c.

In light of eq. (22), we formulate the Foresight-enabled BPR Routing (FBPR, Algorithm 2). The line 5 of the proposed Algorithm reflects the outcome of equation (22). If an alarm level is defined, the search in line 5 is restricted within $c : U_n^{(c)} \geq alarm_level$. The $visited\,[.]$ array is also introduced, to make sure that each possible destination is routed via a single link at each node. The optimization of line 10 acquires a different meaning, pertaining to multi-links. Assume a triple link $M = \{l_1 : \mu_1, l_2 : \mu_2, l_3 : \mu_3\}$ and a corresponding set of $c_l^*(t)$ assignments $A = \left\{ c_{l_1}^*(t), c_{l_2}^*(t), c_{l_3}^*(t) \right\}$. Line 10 refers to the optimal reordering of the assignments out of all possible $M \times A$ combinations and for each multi-link of the network, maximizing the expected throughput. Finally, lines 11-13 install the FBPR-derived priority rules to the corresponding nodes.

Notice that FBPR does not require any critical change over SBPR. To the contrary, taking into account traffic forecasts can be achieved by simply introducing the G terms in line 5. Therefore, FBPR does not alter the general workflow of SBPR and retains the advantage of requiring node-local state information only. In addition, forecast-awareness can be easily made optional by ignoring G in line 5. In this case FBPR falls back to SBPR.

Corollary 1. *FBPR is throughput-optimal.*

We notice that the preceding analysis takes place before the relaxation of equation (8) of the classic analytical procedure. Applying this final relaxation to equation (15) leads to compliance with the Lyapunov stability criterion (relation (3)) and to the proof of throughput optimality, as detailed in [17]. Therefore, FBPR retains this important trait of SBPR as well.

B. On the dependence of the proposed system on the controller.

In order to minimize the network's dependency from the controller and the overhead introduced by the BPR system, the actuation period T should ideally be as large as possible. However, FBPR shows that the routing decisions depend on the prediction of the aggregate traffic that will be generated at each node with the time interval $[t, t+T]$. Furthermore, the accuracy of predictors generally decreases as T increases. Thus, it is important to study the effects of T on the stability of the system.

Returning to inequality (13), we use the theorem of mean value to replace:

$$G_{(n,c)}^{t \to t+T} = \overline{\lambda_{(n,c)}}(t) \cdot T \tag{23}$$

where $\overline{\lambda_{(n,c)}(t)}$ is the average traffic production rate at node n towards c for the time interval $[t, t+T]$. Therefore, we produce:

$$\Delta L(t) \leq T^2 \cdot \sum_{\forall n,c} \left[\mathring{\mu}_{(n)}^{(c)\,2} + \left(\sum_{l:\, d(l)=n} \mu_l^{(c)} + \underset{(n,c)}{\overline{\lambda}}(t) \right)^2 \right]$$
$$- 2T \cdot \sum_{\forall n,c} U_{(n,c)}(t) \cdot \left[\mathring{\mu}_{(n)}^{(c)} - \sum_{l:\, d(l)=n} \mu_l^{(c)} - \underset{(n,c)}{\overline{\lambda}}(t) \right] \tag{24}$$

The RHS of inequality (24) is a quadratic equation of T of the form $RHS_{(24)}(T) = \alpha T^2 - 2\beta T$, with $\alpha > 0$ and roots:

$$\begin{cases} \rho_1 = 0 \\ \rho_2 = \frac{2\beta}{\alpha} \end{cases} \tag{25}$$

Reminding that μ are the nominal capacities of *backbone* links, it is expected that :

$$\left| \frac{2\beta}{\alpha} \right| \approx 0, \text{ for } max\{\mu\} \ggg 1 \text{ bps} \tag{26}$$

Furthermore, it holds that:

$$\frac{d^2 RHS_{(24)}}{d^2 T} = 2\alpha > 0 \tag{27}$$

Thus, the effect of T on the stability of the network can be summarized as follows:

Remark 2. The sensitivity of the bound of the Lyapunov drift in a BPR-based network is proportional to the monitoring and actuation period T.

Lemma 2 states that the increase of T reduces the network stability, for any BPR algorithm (i.e., both SBPR and FBPR). However, the longevity of routing decisions may also be equally influenced by the accuracy of the $\overline{\lambda}_{(n,c)}(t)$ prediction. To demonstrate this fact, we begin by substituting equation (23) in (18):

$$T \cdot \left[\sum_{l:\, d(l)=b(n)} \mu_l^{(c)} + \mathring{\mu}_{(n)}^{(c)} \right] - \left[U_{(n,c)}(t) - \left(U_{(b(n),c)}(t) + \overline{\lambda}_{(b(n),c)}(t) \cdot T \right) \right] = 0,\ \forall n, c \tag{28}$$

Taking the first derivative of the RHS of equation set (28) with regard to $\overline{\lambda}_{(b(n),c)}(t)$ yields:

$$\frac{\partial RHS_{(28)}}{\partial \overline{\lambda}_{(b(n),c)}(t)} = T \tag{29}$$

Reminding that the equation set (28) yields the optimal routing decisions $b(n)$ we deduce that:

Remark 3. The sensitivity of the routing decisions of the FBPR algorithm to the precision of the $\overline{\lambda}_{(n,c)}(t)$ predictions is proportional to T.

Lemmas 2 and 3 both denote the same sensitivity of the system to increases in the actuation period T and the precision of the $\overline{\lambda}_{(b(n),c)}(t)$ prediction. Particularly, the introduction of traffic forecasting-awareness is shown not to introduce a new dominant factor undermining the network stability. Thus, the choice of T can be made jointly for both factors, based on the total effect on the Lyapunov drift as follows:

Definition 4. A set of BP routing decisions is acceptable for a time period T_{max} such that:

$$RHS_{(24)}(T_{max}) \leq L_{max} \tag{30}$$

where L_{max} is a predefined, acceptable constant.

Definition 4 simply ensures that the Lyapunov drift of the network remains in check for a time up to $t + T_{max}$, which takes prediction precision errors into account as well.

C. On the co-operation of BPR and distance-based routing schemes.

The core assumptions of BPR routing regarding the queue dynamics, expressed in eq. (4), allow for the incoming, outgoing and generated node traffic to be attributed to any process. Due to this generality, BPR can operate in parallel with other data transfer systems, such as DVR, without compromising its stability and throughput optimality traits. Nonetheless, the loop-free, latency-aware operation requires further study.

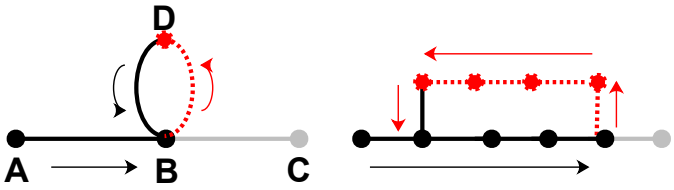

Figure 2: The formation of routing loops is possible when BPR routing and DVR are naively combined. Solid edges correspond to DVR rules, while dashed edges designate BPR rules.

Pure BPR routing (such as SBPR and FBPR) is orthogonal to DVR in terms of its optimization goal. BPR is throughput-oriented and its routing decisions are made on the backpressure potential Δ of equation (21). Reducing the number of hops per routing step is not accounted for. On the other hand, DVR is latency-oriented and its decisions aim at reducing the network's latency. This conflict in operational criteria may lead to the formation of routing loops when DVR and BPR routing decisions are naively combined. Two such cases are illustrated in Fig. 2. At the left inset, data flows are generated from node A towards C. A BPR-derived routing rule intervenes and dictates a detour of the data transfer at node B towards D. However, the DVR policy of D returns the flows back to B creating a loop. The same phenomenon can occur after the intervention of several consecutive BPR rules, as shown at the right inset of Fig. 2.

We proceed to study the conditions that ensure a loop-free combination of BPR and DVR policies. The ensuing formulation will be based on the concept of iterated functions and will refer to the routing of data towards a given destination node c. Let:

$$n' \leftarrow \mathcal{P}(n) \tag{31}$$

denote the application of a routing policy $\mathcal{P}$ on data residing at node n, returning their next intended hop n'. $\mathcal{P}_{BP}$ denotes the use of BPR, while $\mathcal{P}_{DV}$ the use of DVR. In addition, let:

$$n' \leftarrow \mathcal{P}^{\circ m}(n) \tag{32}$$

denote the m-times iterated application of policy $\mathcal{P}$ on node n. Finally, $\overrightarrow{T}\{\mathcal{P}^{\circ m}(n)\}$ will denote the set of traversed nodes according to policy $\mathcal{P}^{\circ m}$.

Lemma 5. *A policy $\mathcal{P}$ is loop-free if-and only if-it has a fixed point on a finite graph.*

Proof: A fixed point of an iterated function $f(.)$ is a point $x : f(x) = x$. Therefore, it holds that $f^{\circ m}(x) = x,\ \forall m$. Assume that a policy $\mathcal{P}$ has a fixed point and a loop originating at a node b, located m steps away from another node n, $b \leftarrow \mathcal{P}^{\circ m}(n)$. Then, for every $m' > m$, $\mathcal{P}^{\circ m'}(n)$ will cycle over the nodes comprising the loop, oscillating around b, without converging. ■

Lemma 6. *The $\mathcal{P}_{BP}$ policy is loop-free.*

Proof: Assume an origin node n_0 and let $\{n_1, n_2, \ldots n_m\} \leftarrow \overrightarrow{T}\{\mathcal{P}^{\circ m}(n_0)\}$. For each step indexed by

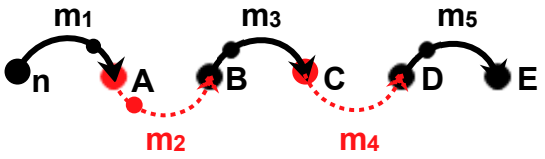

Figure 3: A loop-free combination of $\mathcal{P}_{BP}$ and $\mathcal{P}_{DV}$ policies.

$i = 0 \dots m$, the BPR algorithm has acted based on the Δ metric of equation (21) as follows:

$$\begin{aligned} i=0: &\quad U_{n_0} > U_{n_1} + G_{n_1} \\ i=1: &\quad U_{n_1} > U_{n_2} + G_{n_2} \\ i=2: &\quad U_{n_2} > U_{n_3} + G_{n_3} \\ \dots &\quad \dots \end{aligned}$$

Since the G quantities are strictly positive, it also holds that:

$$U_{n_i} + G_{n_i} > U_{n_{i+1}} + G_{n_{i+1}},\ i = 0 \dots m \tag{33}$$

Therefore, the $\overrightarrow{T}\{\mathcal{P}^{\circ m}(n_0)\}$ set is strictly ordered by descending $U_{n_i} + G_{n_i}$ values. Assume that $\mathcal{P}_{BP}$ creates a loop originating at some step m, which leads back to a node n_j, $j \leq m$ after a number of steps. This would imply that:

$$U_{n_m} + G_{n_m} < U_{n_j} + G_{n_j} \tag{34}$$

which contradicts to equation (33). Thus, $\mathcal{P}_{BP}$ must have a fixed point on a finite graph and, therefore, is loop-free. ■

We proceed to study the conditions for loop-free combinations of $\mathcal{P}_{BP}$ and $\mathcal{P}_{DV}$ policies. Assume that BPR routing has proposed the path illustrated in Fig. 3 for routing data flows from an origin node n towards E. Solid and dashed lines correspond to $\mathcal{P}_{DV}$ and $\mathcal{P}_{BP}$ policies respectively, while m_x denotes the number of nodes comprising each separate piece of the path. First, we note that the path is piece-wise loop-free, since $\mathcal{P}_{BP}$ and $\mathcal{P}_{DV}$ are themselves loop-free. Next, we observe that the path $n \to B$ is loop-free if:

$$\nexists \mathbf{m}:\ A \in \overrightarrow{T}\{\mathcal{P}_{DV}^{\circ \mathbf{m}} \circ \mathcal{P}_{BP}^{\circ m_2} \circ \mathcal{P}_{DV}^{\circ m_1}(n)\} \tag{35}$$

or equivalently:

$$\nexists \mathbf{m}:\ \mathcal{P}_{DV}^{\circ m_1}(n) \in \overrightarrow{T}\{\mathcal{P}_{DV}^{\circ \mathbf{m}} \circ \mathcal{P}_{BP}^{\circ m_2} \circ \mathcal{P}_{DV}^{\circ m_1}(n)\} \tag{36}$$

Setting $\mathcal{P}_{CMB}(n) = \mathcal{P}_{DV}^{\circ m_3} \circ \mathcal{P}_{BP}^{\circ m_2} \circ \mathcal{P}_{DV}^{\circ m_1}(n)$ we proceed to check recursively if the following criteria holds:

$$\nexists \mathbf{m}:\ D \in \overrightarrow{T}\{\mathcal{P}_{DV}^{\circ \mathbf{m}} \circ \mathcal{P}_{BP}^{\circ m_4} \circ \mathcal{P}_{CMB}(n)\},\ \text{or} \tag{37}$$

$$\nexists \mathbf{m}:\ \mathcal{P}_{BP}^{\circ m_4} \circ \mathcal{P}_{CMB}(n) \in \overrightarrow{T}\{\mathcal{P}_{DV}^{\circ \mathbf{m}} \circ \mathcal{P}_{BP}^{\circ m_4} \circ \mathcal{P}_{CMB}(n)\} \tag{38}$$

The process is generalized in the following Lemma:

Lemma 7. *The introduction of a pathlet $\mathcal{P}_{BP}^{\circ m}$ at an intermediate node $\mathcal{P}^*(n)$ of a loop-free path results into a new loop-free path if it holds that:*

$$\nexists \mathcal{M}:\ \mathcal{P}_{BP}^{\circ m} \circ \mathcal{P}^*(n) \in \overrightarrow{T}\left\{\mathcal{P}_{DV}^{\circ \mathcal{M}} \circ \mathcal{P}_{BP}^{\circ m} \circ \mathcal{P}^*(n)\right\} \tag{39}$$

Lemma 7 provides a quick check of loop formation when stitching together $\mathcal{P}_{BP}$ and $\mathcal{P}_{DV}$ policies. Assuming that the topology paths are cache-able, Lemma 7 requires a single linear search per added $\mathcal{P}_{BP}$ pathlet end-point. Furthermore, the check is additive per every new $\mathcal{P}_{BP}$ pathlet. If the check yields loop formation, then BPR is run again at the $\mathcal{P}_{BP}$ pathlet start-point, this time excluding the presently selected neighbor from the search of lines $4-8$, Algorithm 1. If no valid neighbor is found, no BPR priority routing rule is installed at the pathlet start-point. The procedure is formulated as Algorithm 3. The runtime of the algorithm can be limited by an allowed $Timeout$ parameter, expressing a potential deadline imposed by the network regarding the derivation of the priority rules.

Algorithm 3 The FBPR and DVR policy stitching algorithm.

1: **procedure** BP_DV_STITCH($state_{netw}(t), Timeout$)
2: **for each** node n **do**
3: $neighbors[n] \leftarrow \{$nodes $n' : \exists$ link $l_{n \to n'}\}$.
4: **end for**
5: $Loop_found \leftarrow true$.
6: **while** $Loop_found$ **and** $Timeout$ not exceeded,
7: $Loop_found \leftarrow false$.
8: Execute $FBPR(state_{netw}(t), neighbors)$.
9: Detect loops via Lemma 7.
10: **for each** loop-inducing BP pathlet p **do**
11: $Loop_found \leftarrow true$.
12: $neighbors[startpoint\{p\}] -= current_neighbor$.
13: **end for**
14: **end while**
15: Deploy loop-free BP priority rules.
16: **end procedure**

Algorithm 4 Less exploratory FBPR and DVR policy stitching.

1: **procedure** NHOPS_STITCH($state_{netw}(t)$)
2: **for each** origin node n **do**
3: **for each** destination node $c \neq n$ **do**
4: $n' \leftarrow \{n : \left\|\overrightarrow{T}_{(c)}\{\mathcal{P}_{DV}^{\circ\infty}(n)\}\right\| > \left\|\overrightarrow{T}_{(c)}\{\mathcal{P}_{DV}^{\circ\infty}(n')\}\right\|\}$.
5: $neighbors[n][c] \leftarrow n'$.
6: **end for**
7: **end for**
8: Execute $FBPR(state_{netw}(t), neighbors)$.
9: **end procedure**

The BP_DV_STITCH algorithm is intended to promote the exploration of all possible neighbor choices that may offer an alternative exit to locally accumulated traffic via BPR. This exploratory nature enforces the loop control of line 9, as well as the iteration of lines $6-14$. The exploration depth of the algorithm can also be limited deterministically. This choice can offer immunity against issues of asynchronous installation of flow rules, which presently constitutes a significant limitation of OpenFlow-based solutions [40]–[43]. This approach is expressed via Algorithm 4.

Inspired by the study of Ying et al. [20], Algorithm 4, at line 4, limits the BPR search space of neighbors only to those that offer a decreased number of hops towards the destination c. This approach disallows the formation of routing loops, as expressed by the following Lemma:

Lemma 8. *NHOPS_Stitch is loop-free.*

Proof: Consider any path of arbitrary length m from a node n to $c \neq n$, $\{n_1, n_2, \dots, n_m\} = \overrightarrow{T}_{(c)}\{\mathcal{P}^{\circ m}(n)\}$, where

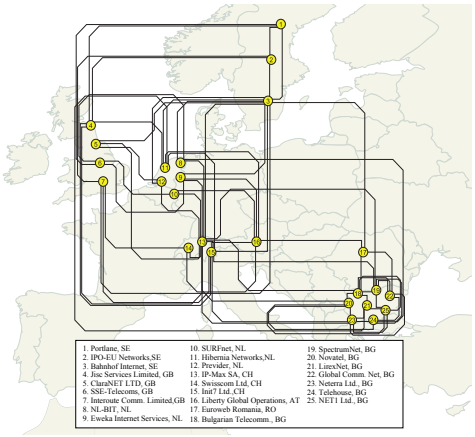

Figure 4: The AS-level topology employed in the simulations, comprising 25 ASes and 66 peering relations. Each node represents an AS and is annotated with its name and location (country code within Europe).

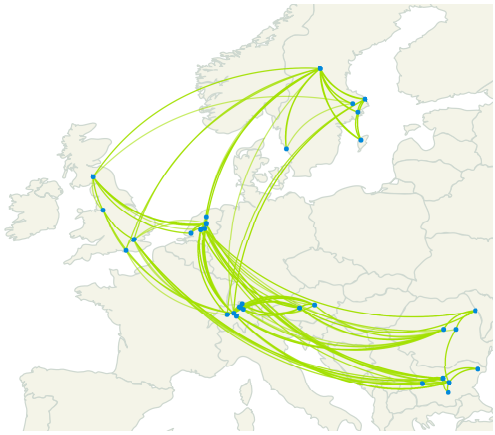

Figure 5: The router-level topology corresponding to Fig 4, containing 351 routers and 273 peering links. (Several router geo-locations coincide at city-level granularity).

$\mathcal{P}^{\circ m}$ comprises any combination of $\mathcal{P}_{BP}$ and $\mathcal{P}_{DV}$ policies. Due to line 4 of Algorithm 4, the nodes $\{n_1, n_2, \ldots, n_m\}$ are order by descending number of hops towards the destination c. As in Lemma 6, the existence of a loop would violate this ordering, concluding the proof. ■

Finally, it holds that:

Lemma 9. *NHOPS_Stitch requires no mechanism to guard against asynchronous or partial deployment of the BP priority rules in the network.*

Proof: The claim is proved by induction. Assume a loop-free path of arbitrary length m from a node n to $c \neq n$, $\{n_1, n_2, \ldots, n_m\} = \overrightarrow{T}_{(c)}\{\mathcal{P}^{\circ m}(n)\}$, where $\mathcal{P}^{\circ m}$ comprises any combination of $\mathcal{P}_{BP}$ and $\mathcal{P}_{DV}$ policies. We install an additional priority flow rule on any of the nodes n_i that are governed by $\mathcal{P}_{DV}$ flow entries only. The resulting path $\overrightarrow{T}'_{(c)}\{\mathcal{P}^{\circ m}(n)\}$ is still loop-free due to Lemma 8. However, there certainly exists at least one loop-free path, namely the $\overrightarrow{T}_{(c)}\{\mathcal{P}^{\circ m}_{DV}(n)\}$, thus proving the claim. ■

V. Simulations

In this Section, the performance of the proposed scheme is evaluated in terms of achieved average throughput, latency and traffic overflow rate in a variety of settings. The employed simulator (implemented on the AnyLogic platform (JAVA) [44]) and datasets are freely available.

The simulations evaluate the potential of collaboration among a cluster of existing ASes. To this end, a real topology of ASes covering the European continent is derived, using the Macroscopic Internet Topology Data Kit (ITDK) by the Center for Applied Internet Data Analysis (CAIDA) [29]. ITDK offers a router-level topology of the Internet. Each router is accompanied by its connectivity (links), its geo-location (city-level) and its AS assignment. From this complete dataset, we derive a simulation-wise tractable subset of 25 ASes. This upper bound was selected via runtime trials in the simulation platform. The studied subset is derived as follows.

Initially, all routers in the ITDK set are mapped to their corresponding ASes. Subsequently, routers that belong to non-Transit ASes are filtered out, using the AS classification dataset by CAIDA [45]. Routers assigned to Tier-1 ASes are discarded as well, since collaboration between such organizations is not straightforward [46]. The resulting router/AS entries are filtered based on their ITDK geolocation, focusing on Europe. Moreover, we retain routers/links that are parts of AS peer-to-peer relations only, using the AS relationships dataset by CAIDA [47]. We sort the remaining ASes by connectivity degree and keep the 25 top-connected ones. The final AS-level topology is illustrated in Fig. 4, accompanied by the AS company names, derived via a CIDR report [48]. The corresponding router-level topology is illustrated in Fig. 5.

It is noted that the employed datasets may contain minor inconsistencies, as reported by the respective publishers. However, the quality of the employed datasets is considered adequate for the scope of the evaluation, which focuses on the efficiency of the proposed inter-AS routing scheme and not on contributing Internet measurements. Additionally, assumptions are needed in the place of certain attributes that are not publicly available. First, the physical or virtual (e.g., remote peering) nature of the links is omitted in the used datasets. Nonetheless, the proposed scheme can operate on any link, regardless of its nature, provided that an AS can freely deploy priority routing rules that refer to it. Second, the ITDK router topology refers to AS border routers, but not to their internal connectivity. Therefore, using a common assumption [49], we consider full-mesh internal AS connectivity. Third, the latency and capacity of Internet links are generally not publicly available. Thus, we assume that the latency of a link is derived by the distance between its endpoints, divided by the speed of light ($3 \cdot 10^8$ m/s). Finally, noticing that 10 Gbps rates are commonly supported by CISCO border routers [50], we pick the capacity of each link at random (uniformly) within the range $[5, 15]$ Gbps at each direction. Multiple runs with different link capacities are executed.

The Controller is placed at the Swiss AS "#13" in Fig. 4, which yields the highest connectivity degree in the examined AS graph. AS routers nearest to the Controller in terms of hops act as NCS endpoints. At the simulation initialization stage,

an instance of the Bellman-Ford algorithm runs at each router, deriving the underlying DVR routing rules. Subsequently, the Controller begins to interact with the NCSes with a period of T sec. All installed priority rules are retained for a full period. Notice that control messages (i.e., affecting the communication of the Controller and the NCSes) receive top routing priority.

Regarding the internal architecture of a router, the endpoints of each router link are connected to router network interfaces (NICs). The aggregate incoming traffic at each router (from all NICs) is first enqueued at a central, shared memory (4 GB [50]), and is subsequently dispatched to the appropriate exit NIC based on the active routing rules. Each NIC is equipped with a 50 MB-sized twin-buffer to avoid idle intervals.

The connection of each border router to the cluster-external Internet is represented by a dedicated network interface. Each router hosts one such NIC, which is connected to an instance of an inter-domain traffic generator (ITMGen) [51]. All such instances within the same AS are identical, and their cumulatively produced traffic rate complies with ITMGen, which specifies a traffic matrix describing the average traffic flow between any AS pair. ITMGen requires as input a metric of mutual popularity, p_{ij}, between any two ASes i, j, which describes the portion of traffic that enters i (cluster-external), destined towards j. In absence of publicly available data, we derive p_{ij} from the connectivity degree d_i of the ASes in the setup of Fig. 4. In other words, we assume that the connectivity degree of an AS reflects its popularity in terms of serving as traffic endpoint. Initially, the aggregate popularity p_i of an AS i is derived as $p_i = \frac{d_i}{\sum_{\forall i} d_i}$. Subsequently, p_{ij} is approximated as $p_{ij} = \frac{p_j}{1-p_i}$, $i \neq j$ and $p_{ii} = 0$. All other configuration parameters of ITMGen are retained [51]. Given that packet-level simulation of backbone networks is generally not tractable in terms of simulation runtimes [5], [52], we assume that the generated traffic is organized in 50 MB-sized batches (i.e., equal to the NIC twin-buffer size).

The logged metrics include the network-wide throughput, overflow rate and average latency. Throughput is calculated as the total traffic volume that has traversed the network links during the simulation, divided by the simulation duration. The overflow metric is defined in a similar fashion, while the latency is measured as the average delivery time of packet-batches. Finally, in the ensuing Figures each separate plot is distinguished by the employed BRP variant (FBPR or SBPR), the DVR-BPR co-existence approach (STITCH for Algorithm 3 and NHOPS for Algorithm 4). FBPR incorporates a simple averaging forecasting method with window size equal to T. Each simulation lasts for 1 hour, which was observed to be sufficient for more than 95% confidence in the logged metrics.

A. Results

Figure 6 evaluates the traffic volume increase potential of the proposed scheme. To this end, the actuation period, T, is set to a value of 10 sec, given that the AS-state monitoring itself requires $3-5$ sec in a real system [28]. The ITMGEN traffic matrix is scaled linearly, creating the load axis values. Boundary and average values over 100 link capacity randomizations are shown.

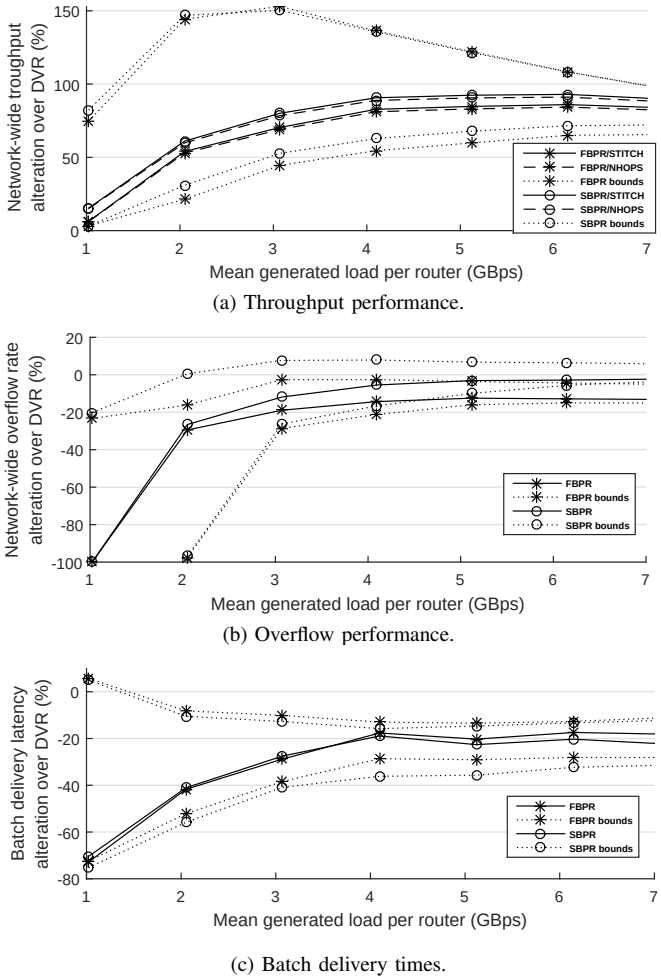

(a) Throughput performance.

(b) Overflow performance.

(c) Batch delivery times.

Figure 6: Evaluation of the proposed scheme versus plain DVR ($T = 10$sec).

According to Fig. 6a, the proposed scheme can increase the average traffic volume by 100% over the plain DVR case. The lower bound is close to the average values, while the upper bound reaches up to 150%. The STITCH-based variations combinations achieve slightly better performance than the NHOPS-based ones. This outcome is expected, given that the NHOPS approach restricts the AS neighborhood search within decreasing hop distances to the end destinations. Naturally, the possible choices for traffic offloading are reduced, which affects the total throughput. Nonetheless, NHOPS retains a high throughput performance, coupled with natural loop-free operation benefits.

The stability of the collaboration cluster also benefits from the proposed scheme, as shown in Fig. 6b. All proposed variations achieve lower average overflow rates than plain DVR. The FBPR variations achieve better results than SBPR, with regard to both average and boundary values. Notice that SBPR yielded marginally better throughput performance in Fig. 6a. Nonetheless, this surplus is overflown, since it can be directed to rapidly overloading ASes. FBPR achieves a better management of traffic by considering future traffic levels.

The batch latency also benefits from the proposed scheme,

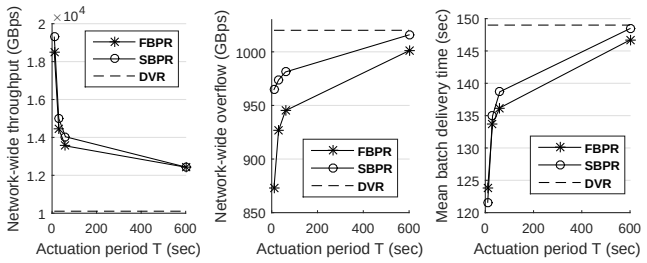

Figure 7: Performance effect of the period T.

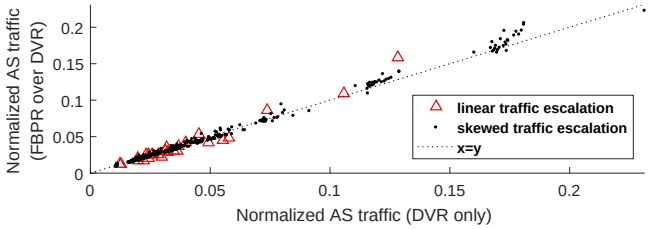

Figure 8: Distribution of surplus traffic to each AS, with and without FBPR.

as shown in Fig. 6c. This observation holds for all examined scheme variations. BPR is known to minimize the amount of queued traffic by distributing it fairly across all collaborating nodes. As a result, the queuing time decreases as well.

The ensuing experiments assume the NHOPS variation only (STITCH is similar) and a medium router load of 4 GBps.

We proceed to study the effects of the actuation period on the performance of the proposed scheme. Figure 7 presents absolute performance results for periods ranging from 10 sec to 10 min. The achieved throughput (left inset) is consistently better than plain DVR, even for the maximal T value. The throughput naturally decreases as T increases, given that the derived priority rules will generally lose their timeliness. The overflow and latency performance is given in the middle and right insets respectively. The relative ranking of the compared schemes is retained. Increasing the actuation period limits the stability and latency benefits of the proposed scheme as well, retaining an improvement over DVR nonetheless.

Given that FBPR and SBPR are both throughput-optimal, their close performance in Fig. 7 (left) is expected. However, in terms of stability (middle), the gains of FBPR are clear (followed by its latency performance - right). Notice that FBPR at, e.g., $T = 600$ sec behaves as SBPR at $T = 350$ sec. Thus, ASes can also use forecasting for increasing their Controller interaction period, while retaining SBPR-level stability.

We proceed to study the distribution of the transit traffic over the collaborating ASes. We perform runs with different cases of router load escalation and log the throughput traffic percentage per AS under: i) plain DVR and ii) FBPR/NHOPS/T=10sec over DVR. The *linear* escalation case is the proportional escalation of the ITMGEN traffic matrix leading to 4 GBps average router load. In the *skewed* escalation case, the same traffic load enters the network, but solely from one given AS. In the skewed case, 25 separate runs are executed, with one AS acting as the traffic input point. The ensuing throughput traffic distribution assumes the non-input ASes. The DVR and FBPR/DVR results are given as a scatter plot in Fig. 8. All skewed-case runs are collectively plotted.

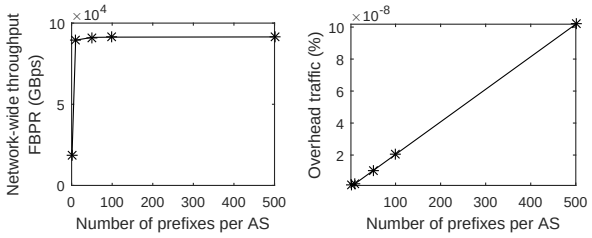

Figure 9: FBPR operation at prefix-level and incurred overhead.

Figure 8 exhibits a strongly linear relation between the traffic distribution over plain DVR, and the corresponding distribution when FBPR is activated. Thus, the traffic increase due to FBPR at each AS is fair and proportional to its original DVR traffic, in accordance with BPR theory (cf. end of Section II). In other words, FBPR will generally not alter the original AS ranking in terms of served traffic volume. The SBPR behavior is similar and is omitted for Figure clarity.

So far, the evaluation assumed traffic routing at AS-granularity. Figure 9 studies the effect of routing at prefix-granularity on the proposed scheme. We focus on FBPR (SBPR is similar), using the default mean router load (4 GBps) and $T = 10$ sec. Each AS hosts a varying number of prefixes N_P (x-axis of Fig. 9), and uses separate backlogs ($U(t)$) for each one. The FBPR operation follows the steps of Algorithm 2, with the modification that N_P/N_L priority rules can be assigned to each router link, N_L being the number of outgoing router links. Moreover, the network overhead imposed by FBPR is logged, assuming the worst case scenario where each AS reports its full internal congestion per prefix every $T = 10$ sec. The report is a set comprising the AS identifier (once, 8-byte string), prefix IPs (4 bytes each), prefix subnets (1 byte each) and prefix load in GB (8-byte float each). A priority rule proposal is a tuple comprising the intended AS identifier (once), and the prefix IPs (4 bytes), subnets (1 bytes) and link identifiers (8-byte string).

As shown in Fig. 9 (left inset), per-prefix routing naturally increases the achieved network throughput, since it allows for more fine-grained traffic management compared to per-AS routing. Even for a small number of prefixes ($N_p = 10$), the performance of FBPR increases by a factor of $\times 5$, essentially reaching the maximal value. The imposed overhead is trivial (Fig. 9 - right inset), amounting at $10^{-8}\%$ of the overall traffic volume, while scaling linearly with the number of prefixes.

We note that, in Fig. 9, all ASes are assumed to manage network prefixes in the same manner. In reality, an AS_1 may treat several prefixes individually, while a neighbor AS_2 may treat them as a single super-prefix. Both SBPR and FBPR can operate without changes in the AS prefix management. When considering backlog differences, SBPR and FBPR simply then consider the affected backlogs at each side. E.g., offloading AS_2 to AS_1 will consider the backlog difference of the load for the super-prefix at AS_2 minus the cumulative load of all affected sub-prefixes at AS_1.

Finally, Fig. 10 studies the effect of topology size on the performance of the proposed scheme. We assume the default

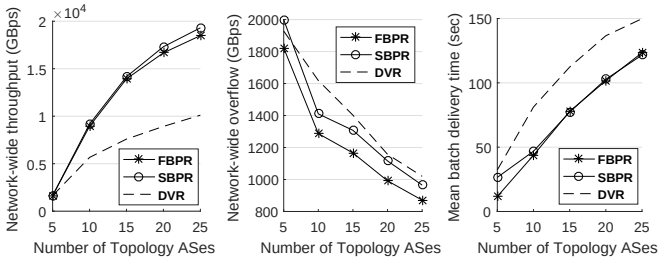

Figure 10: Performance effect of the topology size.

router load (4 GBps), T =10 sec and per-AS routing. We select size-varying subsets of the original 25 ASes (producing the corresponding router-level topologies as well), forming the x-axis of Fig. 10. Link capacities are randomized 100 times, presenting average performance values in the y-axis.

Naturally, vary small topologies (5 ASes) offer limited path diversity and network capacity. The paths are shorter nonetheless, yielding decreased latency. However, at slightly bigger topologies (10 ASes), the throughput, stability and latency benefits of backpressure appear. The improvement over plain DVR then increases with the topology size, as new paths are added and exploited by the BPR schemes.

B. Discussion and Future Steps

The evaluation focused on factors that can constitute the proposed scheme appealing for initiating inter-AS collaboration. Specifically, three benefits were highlighted: i) Attaining a significant increase in the transit traffic volume handled by the AS-cluster. ii) A natural gain-sharing mechanism among the collaborating ASes. Each particular AS sees an increase in its own transit traffic based on its centrality within the AS-topology. iii) Lightweight implementation, requiring a web service accessible by the ASes. The point-of-failure risk is expected to be low, given that the scheme then naturally falls back to standard AS routing. Emphasis is placed on using the NCSes that ASes have already adopted. This means that the proposed routing rules are not deployed automatically by the proposed scheme, but rather passed to each NCS for further processing without overriding it. This approach allows the ASes to retain the full control of their networks.

Nonetheless, the proposed scheme assumes ASes willing to cooperate. Trust and information exchange are sensitive matters in commercial relations. The proposed scheme requires ASes to trust it with partial information on their internal congestion state. For instance, an AS_1 may inform the proposed scheme that "10 GB are pending towards prefix p". This information exchange requires a degree of trust. Nonetheless, ASes can already obtain similar information via measurements. For example, an AS_2 that neighbors AS_1 can easily derive that "10 GB were received from AS_1 towards prefix p" and yield a forecast on the present state of AS_1. Thus, the present scheme mainly requires a change in the way such information is exchanged and not in its confidentiality.

Future work is directed towards the important aspect of security against AS foul play, and particularly against reporting false congestion levels $U_{(n,c)}(t)$. The concept of promise/deliverable can constitute the basis for detecting such attempts. For example, if an AS reports falsely high $U_{(n,c)}(t)$ levels to elicit assistance, it should also deliver this amount of real traffic. This can be verified via measurements by the assisting ASes. In cases where falsely low $U_{(n,c)}(t)$ levels are reported to attract traffic, the attracting AS should yield low latency and loss rate in handling this traffic. Failure to deliver can imply foul-play, evicting untrustworthy ASes from the cluster.

VI. Related Work

Highly-efficient and even optimal TE within datacenters constituted one of the earliest successes of SDN. Emphasis was placed on versatile traffic prioritization systems at flow granularity, as well as network-wide optimization objectives such as throughput maximization. B4 incorporates this concern by keeping tuples of source, destination and QoS traits per network flow [6]. The network's resources are constantly monitored and the flows are assigned paths according to their priority. B4 is also known for achieving near-optimal network throughput. Microsoft's SWAN considers classes of priorities, pertaining to critical, elastic and background traffic [2]. Network paths are first assigned per priority class. Within each coarse assignment, a max-min fairness approach is used to distribute resources to specific flows. Bell Labs propose a more direct approach, seeking to solve the formal maximal link utilization problem, given explicit flow requests [3]. Other studies focus on scenarios such as partially SDN-controlled networks, or on multipath routing [53], exploiting the monitoring capabilities of OpenFlow [53].

Related studies have also demonstrated the SDN potential in congestion handling, network monitoring and application specific QoS. MicroTE [54], Hedera [52] and Mahout [5] focus on the detection and special handling of large "elephant" flows, under the assumption that they constitute the usual suspects of congestion. Such flows are assigned to paths which do not conflict with the bulk of the remaining traffic. The network monitoring is continuous, scanning network-wide for large flows via periodic polling at the scale of 5 sec. A similar high-level logic is adopted for application-specific TE. Initially, applications state their latency and bandwidth requirements. Then, the corresponding flows are mapped to appropriate paths. PlugNserve assumes application-specific TE and focuses on robustness against adding/removing nodes in real-time [55]. The Aster*x approach seeks to keep the response time of web services as low as possible, by monitoring the congestion levels of the network and picking the less loaded routes per new flow [56]. Authors in [57] propose a scheme that assigns flows to routes, selected from a list of paths with similar bandwidth, ordered by ascending latency.

The scalability of SDN-enabled TE has constituted an early consideration. Authors in [1] stress the need to reduce the control-plane load by: (i) minimizing the required amount of flow rules installed to the network, and (ii) limiting the interaction between the SDN controller and the routers. Towards the first direction, Devoflow [58] and DIFANE [59] introduced an operation promoting wild-card rules, in order to minimize the required flow entries. This fact is taken into consideration in routing update processes, where older

routing rules are retained for some time, in order to transit smoothly from one TE instance to another [41]–[43]. Towards the controller workload direction, present solutions multiple controller deployments, with workload balancing mechanisms [59], [60]. Kandoo, for example, proposes a hierarchical controller deployment comprising two layers of command. The lowest layer handles dedicated partitions of the network, while a central controller coordinates any intra-partition action [61].

Noticing the operational benefits and the growing maturity of SDN, researchers studied its fitness as a platform for evolving the inter-AS routing. The ossification of Internet routing, the absence of QoS guarantees and the potentially slow convergence and security issues of BGP constitute some of the long-standing issues [11]. The 4D approach constituted an early *clean-slate* proposal for inter-AS routing [7]. 4D can be considered as the precursor of OpenFlow, and allows AS operators to set clear network objectives, obtain network-wide views and directly control the inter-network state from a central point. Authors in [11] proposed an evolvable platform for inter-AS routing based on the principle of control-plane outsourcing. The routing logic of multiple ASes is outsourced to an external trusted entity, which orchestrates the inter-AS collaborative routing. SDN serves as the underlying technology, due to the clear separation it enforces between the control and data planes. Platforms that treat the routing infrastructure as a service constitute specific, fitting choices (E.g., RouteFlow [62]). In order to limit the scale of the required changes in hardware and protocol for adopting these solutions, studies on partial deployment have taken place [10]. Highly attractive gains in routing convergence are attained when 50% of the ASes convert to the new scheme. The SDX approach proposed the software-ization of Internet Exchanged Points (IXPs), central rendezvous points for AS peering that grow in popularity [9]. The Control Exchange Points approach proposes that ASes publish some internal paths to an external entity [8]. The entity can then stitch together paths crossing multiple ASes, offering end-to-end QoS upon demand.

The present work is based on the premise that ASes may require a minimal-commitment scheme at first, in order to try-out and evaluate the prospects of collaboration. Therefore, we propose a TE approach that can operate over existing hardware and NCSes, without introducing point-of-failure considerations. The proposed scheme constitutes a novel application of BPR routing [15], and it brings its analytically-proven, latency-aware throughput-optimality to inter-AS TE, with limited commitment. The proposed approach prioritizes compatibility with existing AS operations, differing from clean-slate approaches [63]. The proposed system is not intended to offer the rich set of features promised by related studies, but rather to serve as an intermediate step towards convincing ASes to gradually adopt them. The scheme assumes SDN-inspired principles for its operation, seeking to lay the foundations for more extensive SDN inter-AS adoption in the future. ASes that adopt the proposed scheme can benefit from the optimal stability it entails. CXP pathlet stitching can serve as the next step towards closer collaboration, taking advantage of the capabilities offered by SDX points for end-to-end QoS.

An early version of the proposed scheme is given at [64].

VII. Conclusion

The present study proposed the application of BPR routing as a minimal-commitment scheme for collaborative inter-AS traffic engineering. BPR itself promises network throughput maximization and transit traffic increase, which can supply an economic incentive for AS participation to the proposed scheme. Novel BPR-based routing algorithms where analyzed for the inter-AS traffic engineering scenario. From a systemic aspect, BPR is deployed to the network of collaborating ASes via a simple SDN-inspired interface, in the form of temporary priority routing rules. Extensive simulations evaluated the further advantages of the proposed scheme, namely stability under increased network load and co-existence with existing inter-AS routing mechanisms. Serving as a first step towards deeper AS cooperation could constitute a significant application of the proposed approach.

References

[1] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner, "OpenFlow: enabling innovation in campus networks," *ACM SIGCOMM Computer Communication Review*, vol. 38, no. 2, pp. 69–74, 2008.

[2] C.-Y. Hong, S. Kandula, R. Mahajan, M. Zhang, V. Gill, M. Nanduri, and R. Wattenhofer, "Achieving high utilization with software-driven WAN," in *Proceedings of the ACM SIGCOMM conference*, 2013, pp. 15–26.

[3] S. Agarwal, M. Kodialam, and T. V. Lakshman, "Traffic engineering in software defined networks," in *INFOCOM, 2013 Proceedings IEEE*, 2013, pp. 2211–2219.

[4] S. Kandula, D. Katabi, B. Davie, and A. Charny, "Walking the tightrope: Responsive yet stable traffic engineering," in *ACM SIGCOMM Computer Communication Review*, vol. 35, 2005, pp. 253–264.

[5] A. R. Curtis, W. Kim, and P. Yalagandula, "Mahout: Low-overhead datacenter traffic management using end-host-based elephant detection," in *INFOCOM, 2011 Proceedings IEEE*, 2011, pp. 1629–1637.

[6] S. Jain, A. Kumar, S. Mandal, J. Ong, L. Poutievski, A. Singh, S. Venkata, J. Wanderer, J. Zhou, M. Zhu *et al.*, "B4: Experience with a globally-deployed software defined WAN," in *Proceedings of the ACM SIGCOMM conference*, 2013, pp. 3–14.

[7] A. Greenberg, G. Hjalmtysson, D. A. Maltz, A. Myers, J. Rexford, G. Xie, H. Yan, J. Zhan, and H. Zhang, "A clean slate 4d approach to network control and management," *SIGCOMM Comput. Commun. Rev.*, vol. 35, no. 5, pp. 41–54, Oct. 2005.

[8] V. Kotronis, M. Rost, B. A. Georgopoulos, B. Ager, S. Schmid, and X. Dimitropoulos, "Stitching inter-domain paths over ixps," *ACM SOSR*, 2016.

[9] A. Gupta, L. Vanbever, M. Shahbaz, S. P. Donovan, B. Schlinker, N. Feamster, J. Rexford, S. Shenker, R. Clark, and E. Katz-Bassett, "Sdx: A software defined internet exchange," *ACM SIGCOMM Computer Communication Review*, vol. 44, no. 4, pp. 551–562, 2015.

[10] P. Sermpezis and X. Dimitropoulos, "Inter-domain sdn: Analysing the effects of routing centralization on bgp convergence time," in *MAMA workshop (ACM SIGMETRICS)*, 2016.

[11] V. Kotronis, A. Gämperli, and X. Dimitropoulos, "Routing centralization across domains via sdn: A model and emulation framework for bgp evolution," *Computer Networks*, vol. 92, pp. 227–239, 2015.

[12] I. Castro, R. Stanojevic, and S. Gorinsky, "Using Tuangou to Reduce IP Transit Costs," *IEEE/ACM Transactions on Networking*, vol. 22, no. 5, pp. 1415–1428, 2014.

[13] L. Gyarmati, R. Stanojevic, M. Sirivianos, and N. Laoutaris, "Sharing the cost of backbone networks: cui bono?" in *ACM IMC 2012*, p. 509.

[14] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," *Automatic Control, IEEE Transactions on*, vol. 37, no. 12, pp. 1936–1948, 1992.

[15] L. Georgiadis, M. J. Neely, and L. Tassiulas, "Resource Allocation and Cross-Layer Control in Wireless Networks," *FNT in Networking (Foundations and Trends in Networking)*, vol. 1, no. 1, pp. 1–144, 2005.

[16] M. J. Neely, E. Modiano, and C. E. Rohrs, "Power allocation and routing in multibeam satellites with time-varying channels," *IEEE/ACM Transactions on Networking*, vol. 11, no. 1, pp. 138–152, 2003.

[17] ——, "Dynamic power allocation and routing for time-varying wireless networks," *IEEE Journal on Selected Areas in Communications*, vol. 23, no. 1, pp. 89–103, 2005.

[18] E. Leonardi, M. Mellia, F. Neri, and M. Ajmone Marsan, "Bounds on average delays and queue size averages and variances in input-queued cell-based switches," in *IEEE INFOCOM 2001 - IEE Conference on Computer Communications*, 22-26 April 2001, pp. 1095–1103.

[19] N. McKeown, A. Mekkittikul, V. Anantharam, and J. Walrand, "Achieving 100% throughput in an input-queued switch," *IEEE Transactions on Communications*, vol. 47, no. 8, pp. 1260–1267, 1999.

[20] L. Ying, S. Shakkottai, A. Reddy, and S. Liu, "On Combining Shortest-Path and Back-Pressure Routing Over Multihop Wireless Networks," *IEEE/ACM Trans. on Networking*, vol. 19, no. 3, pp. 841–854, 2011.

[21] L. Huang, S. Moeller, M. J. Neely, and B. Krishnamachari, "LIFO-Backpressure Achieves Near-Optimal Utility-Delay Tradeoff," *IEEE/ACM Trans. on Networking*, vol. 21, no. 3, pp. 831–844, 2013.

[22] L. Huang and M. J. Neely, "Delay reduction via Lagrange multipliers in stochastic network optimization," *IEEE Transactions on Automatic Control*, vol. 56, no. 4, pp. 842–857, 2011.

[23] S. Moeller, A. Sridharan, B. Krishnamachari, and O. Gnawali, "Routing without routes," in *the 9th ACM/IEEE International Conference*, T. Abdelzaher, T. Voigt, and A. Wolisz, Eds., p. 279.

[24] H. Seferoglu and E. Modiano, "TCP-aware backpressure routing and scheduling," in *2014 Information Theory and Applications Workshop (ITA)*, pp. 1–9.

[25] P. Berde, M. Gerola, J. Hart, Y. Higuchi, M. Kobayashi, T. Koide, B. Lantz, B. O'Connor, P. Radoslavov, W. Snow *et al.*, "Onos: towards an open, distributed sdn os," in *Proceedings of the third workshop on Hot topics in software defined networking*. ACM, 2014, pp. 1–6.

[26] S. H. Yeganeh and Y. Ganjali, "Beehive: Simple distributed programming in software-defined networks," in *Proceedings of ACM SOSR'16*. ACM, pp. 4:1–4:12.

[27] O. Akashi, K. Fukuda, T. Hirotsu, and T. Sugawara, "Policy-based BGP-control architecture for inter-AS routing adjustment," *Computer Communications*, vol. 31, no. 13, pp. 2996–3002, 2008.

[28] A. Tootoonchian, M. Ghobadi, and Y. Ganjali, "OpenTM: traffic matrix estimator for OpenFlow networks," in *Passive and Active Measurement*, 2010, pp. 201–210.

[29] Center for Applied Internet Data Analysis, "Macroscopic Internet Topology Data Kit (ITDK)," 2016. [Online]. Available: https://www.caida.org/data/internet-topology-data-kit

[30] L. Vanbever, S. Vissicchio, C. Pelsser, P. Francois, and O. Bonaventure, "Seamless network-wide IGP migrations," *ACM SIGCOMM Computer Communication Review*, vol. 41, no. 4, pp. 314–325, 2011.

[31] T. Thapngam, S. Yu, W. Zhou, and G. Beliakov, "Discriminating DDoS attack traffic from flash crowd through packet arrival patterns," in *IEEE INFOCOM 2011 - IEEE Conference on Computer Communications Workshops*, pp. 952–957.

[32] M. S. Kang, S. B. Lee, and V. D. Gligor, "The crossfire attack," in *Proc. of Security and Privacy (SP)*, 2013, pp. 127–141.

[33] A. Studer and A. Perrig, "The Coremelt Attack," in *Proceedings of the 14th European Conference on Research in Computer Security*, ser. ESORICS, 2009.

[34] M. Prince, "The DDoS That Almost Broke The Internet," *[Online] https://blog.cloudflare.com/the-ddos-that-almost-broke-the-internet/*, 2013.

[35] S. B. Lee, M. S. Kang, and V. D. Gligor, "CoDef: Collaborative Defense Against Large-scale Link-flooding Attacks," in *Proceedings of the Ninth ACM Conference on Emerging Networking Experiments and Technologies*, ser. CoNEXT '13, 2013.

[36] T. Otoshi, Y. Ohsita, M. Murata, Y. Takahashi, K. Ishibashi, and K. Shiomoto, "Traffic prediction for dynamic traffic engineering," *Computer Networks*, vol. 85, pp. 36–50, 2015.

[37] A. Steland, E. Rafajłowicz, and K. Szajowski, Eds., *Stochastic Models, Statistics and Their Applications*, ser. Springer Proceedings in Mathematics & Statistics. Cham: Springer International Publishing, 2015.

[38] P. Cortez, M. Rio, M. Rocha, and P. Sousa, "Multi-scale Internet traffic forecasting using neural networks and time series methods," *Expert Systems*, vol. 29, pp. 143–155, 2010.

[39] J.-B. Hiriart-Urruty, J.-J. Strodiot, and V. H. Nguyen, "Generalized Hessian matrix and second-order optimality conditions for problems withC 1,1 data," *Applied Mathematics & Optimization*, vol. 11, no. 1, pp. 43–56, 1984.

[40] M. Reitblatt, N. Foster, J. Rexford, and D. Walker, "Consistent updates for software-defined networks: Change you can believe in!" in *Proceedings of the 10th ACM Workshop on Hot Topics in Networks*, 2011, p. 7.

[41] M. Reitblatt, N. Foster, J. Rexford, C. Schlesinger, and D. Walker, "Abstractions for network update," in *Proceedings of the ACM SIGCOMM conference*, 2012, pp. 323–334.

[42] R. McGeer, "A safe, efficient update protocol for OpenFlow networks," in *Proceedings of the first workshop on Hot topics in software defined networks*, 2012, pp. 61–66.

[43] N. P. Katta, J. Rexford, and D. Walker, "Incremental consistent updates," in *Proceedings of the second ACM SIGCOMM workshop on Hot topics in software defined networking*, 2013, pp. 49–54.

[44] X. Technologies, "The AnyLogic Simulator," 2013. [Online]. Available: http://www.xjtek.com/anylogic/

[45] Center for Applied Internet Data Analysis, "AS Classification Dataset," 2016. [Online]. Available: https://www.caida.org/data/as-classification/

[46] Dyn Research (Renesys), "Annual AS Rankings," 2015. [Online]. Available: http://research.dyn.com/2016/04/a-bakers-dozen-2015-edition/

[47] Center for Applied Internet Data Analysis, "AS Relationships," 2016. [Online]. Available: http://www.caida.org/data/as-relationships/

[48] T. Bates, P. Smith, and G. Huston, "CIDR Report for 1-Mar-2016," 2016. [Online]. Available: http://www.cidr-report.org/as2.0/autnums.html

[49] T. Bates, E. Chen, and R. Chandra, "BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)," *IETF RFC 4456*, 2006. [Online]. Available: http://www.rfc-editor.org/rfc/rfc4456.txt

[50] CISCO, "ASR 9000 Series Routers," 2015. [Online]. Available: http://www.cisco.com/c/en/us/products/collateral/routers/asr-9000-series-aggregation-services-routers/data_sheet_c78-712040.html

[51] J. Mikians, N. Laoutaris, A. Dhamdhere, and P. Barlet-Ros, "ITMgen: A first-principles approach to generating synthetic interdomain traffic matrices," in *Proceedings of the IEEE International Conference on Communications (ICC)*, 2013, pp. 2507–2512.

[52] M. Al-Fares, S. Radhakrishnan, B. Raghavan, N. Huang, and A. Vahdat, "Hedera: Dynamic Flow Scheduling for Data Center Networks," in *Proceedings of the 7th USENIX Conference on Networked Systems Design and Implementation*, ser. NSDI'10. Berkeley, CA and USA: USENIX Association, 2010.

[53] J. Domżał, Z. Duliński, M. Kantor, J. Rząsa, R. Stankiewicz, K. Wajda, and R. Wójcik, "A survey on methods to provide multipath transmission in wired packet networks," *Comp. Networks*, vol. 77, pp. 18–41, 2015.

[54] T. Benson, A. Anand, A. Akella, and M. Zhang, "Microte: fine grained traffic engineering for data centers," in *Proceedings of the seventh CONEXT conference*, 2011, p. 8.

[55] N. Handigol, S. Seetharaman, M. Flajslik, N. McKeown, and R. Johari, "Plug-n-Serve: Load-balancing web traffic using OpenFlow," *ACM SIGCOMM Demo*, 2009.

[56] N. Handigol *et al.*, "Aster*x: Load-Balancing Web Traffic over Wide-Area Networks," *Stanford University, Online Article http://yuba.stanford.edu/ nikhilh/pubs/handigol-gec9.pdf*, 2014.

[57] N. Gvozdiev, B. Karp, and M. Handley, "FUBAR: Flow Utility Based Routing," in *Proceedings of the 13th ACM Workshop on Hot Topics in Networks*, 2014, p. 12.

[58] J. C. Mogul, J. Tourrilhes, P. Yalagandula, P. Sharma, A. R. Curtis, and S. Banerjee, "Devoflow: Cost-effective flow management for high performance enterprise networks," in *Proceedings of the 9th ACM SIGCOMM Workshop on Hot Topics in Networks*, 2010, p. 1.

[59] M. Yu, J. Rexford, M. J. Freedman, and J. Wang, "Scalable flow-based networking with DIFANE," *SIGCOMM Comput. Commun. Rev.*, vol. 41, no. 4, pp. –, 2010.

[60] Y. Hu, W. Wang, X. Gong, X. Que, and S. Cheng, "BalanceFlow: Controller load balancing for OpenFlow networks," in *Cloud Computing and Intelligent Systems (IEE CCIS)*, vol. 2, 2012, pp. 780–785.

[61] S. Hassas Yeganeh and Y. Ganjali, "Kandoo: a framework for efficient and scalable offloading of control applications," in *Proceedings of the first workshop on Hot topics in SDNs*, 2012, pp. 19–24.

[62] M. Nascimento *et al.*, "Virtual routers as a service: the routeflow approach leveraging software-defined networks," in *the ACM International Conference on Future Internet Technologies 2011*, pp. 34–37.

[63] L. Ying, R. Srikant, and D. Towsley, "Cluster-based back-pressure routing algorithm," in *IEEE INFOCOM 2008*, 2008, pp. 484–492.

[64] C. Liaskos, "A Lightweight, Non-intrusive Approach for Orchestrating Autonomously-managed Network Elements," in *IEEE ISCC'15*.

Christos Liaskos received the Diploma in Electrical and Computer Engineering from the Aristotle University of Thessaloniki (AUTH), Greece in 2004, the MSc degree in Medical Informatics in 2008 from the Medical School, AUTH and the Ph.D. degree in Computer Networking from the Dept. of Informatics, AUTH in 2014. He has published work in venues such as IEEE Transactions on: Networking, Computers, Vehicular Technology, Broadcasting, Systems Man & Cybernetics, Communications, IEEE INFOCOM, Elsevier NANOCOMNET, ACM NANOCOM, at a total of more than 40 publications, and has received a best paper award. He is currently a researcher at the Foundation of Research and Technology, Hellas (FORTH). His research interest include traffic engineering and security in software-defined networks, wireless networks and nanonetworks.

Xenofontas Dimitropoulos is an Assistant Professor at the Institute of Computer Science Department, University of Crete (UoC) since January 2014. In addition, he is affiliated with the Institute of Computer Science of the Foundation for Research and Technology Hellas (FORTH). He leads a research group working on Internet measurements and software defined networks with the main goal of making the Internet more reliable and secure. Before, he worked in the Georgia Institute of Technology, the IBM Research Labs and the University of California San Diego (UCSD). He has published more than 88 papers and 3 patents and has received 2 best paper awards. In addition, he has received prestigious grants from the European Research Council, the Marie Curie and the Fulbright programs. He has served in the Organizing and Technical Program Committee of the flagship networking conference, ACM SIGCOMM.

Leandros Tassiulas (S'89–M'91–SM'05–F'07) received the Ph.D. degree in electrical engineering from the University of Maryland, College Park, MD, USA, in 1991. He has been a Faculty Member with the NYU Polytechnic School of Engineering, Brooklyn, NY, USA, the University of Maryland, College Park, and the University of Thessaly, Volos, Greece. He is currently the John C. Malone Professor of Electrical Engineering with Yale University, New Haven, CT, USA. His most notable contributions include the max-weight scheduling algorithm and the back-pressure network control policy, opportunistic scheduling in wireless, the maximum lifetime approach for wireless network energy management, and the consideration of joint access control and antenna transmission management in multiple antenna wireless systems. His research interests include computer and communication networks with an emphasis on fundamental mathematical models and algorithms of complex networks, architectures and protocols of wireless systems, sensor networks, novel internet architectures, and experimental platforms for network research. He was a recipient of several awards, including the IEEE Koji Kobayashi Computer and Communications Award, the Inaugural INFOCOM 2007 Achievement Award for fundamental contributions to resource allocation in communication networks, the INFOCOM 1994 Best Paper Award, the National Science Foundation (NSF) Research Initiation Award (1992), the NSF CAREER Award (1995), the Office of Naval Research Young Investigator Award (1997), and the Bodossaki Foundation Award (1999).